

October 1982

Dynamic-RAM Controller Orchestrates Memory Systems

Jim Nadir
Mel Bazes
Intel Corporation
Santa Clara, California

"Reprinted from ELECTRONICS, Sept. 8, 1982. Copyright (c) McGraw-Hill, Inc. 1982. All rights reserved."

OCTOBER 1982
ORDER NUMBER: 210758-001

Dynamic-RAM controller orchestrates memory systems

Up to 88 chips take their cues from an n-channel MOS IC that both housekeeps and supports error-corrected dual-port memories

by Jim Nadir and Mel Bazes, Intel Corp., Santa Clara, Calif.

□ Designing a dynamic-random-access-memory system means balancing the goals of high performance, reliability, and versatility against the often contrary aims of economy, simplicity, and compactness. In the last five or so years, the advent of dynamic-RAM controller chips relieved designers of some of the onus of tending to the needs of dynamic chips: standard supportive integrated circuits brought together the counters, timers, multiplexers, and other elements needed.

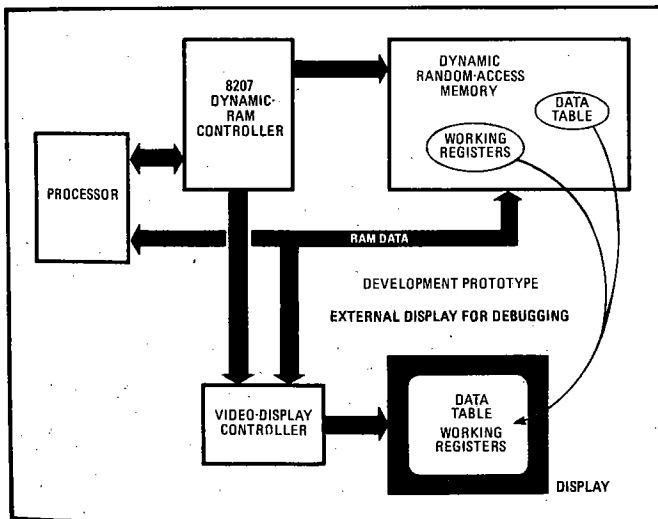
But controllers diverged into two types. One bought the high performance to ride with fast memory systems at the expense of functionality, while the other took on more and more functions to do a complete but slower job. The 8207—an advanced dynamic-RAM controller—blunts the horns of that dilemma and also solves a variety of less severe design problems.

A dynamic-RAM controller is charged with making a dynamic memory system appear static to the host processor. At a minimum, therefore, the controller takes over refreshing the memory chips, multiplexing the row and column addresses, generating control signals, timing the precharge period, and signaling the processor when

data is available or no longer needed. But, beyond those local housekeeping chores, the controller can also go a long way to solving more global design problems, like sharing memory between two processors, not to mention detecting and correcting errors.

To realize this potential for a highly integrated solution, the 8207 has a dual-port interface and, when used with the 8206 error-checking and -correction unit, ensures data integrity in large dynamic-RAM systems. In addition to doing the jobs of refreshing, address multiplexing, and control timing, the unit supports memory-bank interleaving for pipelined accesses, overlaying RAM and read-only-memory locations, and initializing RAM.

The exact implementation of most of these functions is programmable, letting designers tailor their systems in detail. Systems containing up to 88 dynamic-RAM chips—whether 16-, 64-, or 256-K versions—in one, two, or four banks need only a single 8207 and no external buffering. Attesting to the high performance claimed, the 8207 mates dynamic RAMs having 100-nanosecond access times to the iAPX-286 processor operating at 8 megahertz without introducing any wait states.



1. Window on a micro. One use for a dual-port memory shared by independent processors is the development system shown. Adding a video display to the prototype itself gives a window on the system memory.

To achieve that speed and include all those functions, the 8207 relies on a dense, high-speed n-channel MOS process (H-MOS II) and requires a chip some 230 by 200 mils in area. To meet the rigors of operation with even faster processors, novel logic and integrated-circuit designs are employed. Replacing the two-phase logic common in n-MOS ICs, single-phase edge-triggered logic simplifies logic and circuit design, precludes problems of clock-pulse overlap, and reduces the sensitivity to clock high and low times. Voltage-controlled capacitive loads form the delay elements that time critical output pulses, such as the address strobes, and compensate the output-switching delays for variations in power-supply voltage, temperature, and processing.

A low 20-ns setup time for input signals is achieved by cutting the RC delay of input-protection devices and moving the TTL-to-MOS signal buffering from the input pads to the pulse generators. A short 35-ns delay from input to output switching is achieved by triggering the output generators directly from the external clock, saving a buffer delay time. With the resulting high-speed performance and a high level of integration, the 8207 successfully attacks the stringent requirements of today's memory systems.

One system feature gaining popularity currently is the use of multiple processors operating on shared data to obtain higher performances and reliability. For example, a separate processor dedicated to input/output tasks frees the main processor for full-time data processing. Alternatively, multiple main processors can execute different tasks simultaneously. In all such cases, sharing a common memory space among the cooperating processors is the key to effective operation.

Unfortunately, when more than one processor accesses shared memory through a single bus, the limited bus bandwidth and the time spent in exchanging bus control slow down data transfers. Dual-port memory systems overcome this limitation by giving two processors access

to a common memory through two independent buses. The 8207 includes a dual-port interface to simplify the design of shared memory systems.

Two-port memories can be used with multiprocessing or multitasking architectures. In the former, independent processors run independent programs, sharing only a common memory. Multitasking processors cooperate on different parts of the same task.

An example of a multiprocessing architecture is the dynamic video display (Fig. 1) that provides a window on a processor's memory. Centering the display over a data table, for example, immediately reveals how program execution affects the data, which aids in debugging programs. If a microcomputer is implemented with a dual-port memory—the second port for a dynamic video display—then the prototype itself can serve as a development and debugging system, reverting to single-port operation in the final version.

A dual-port architecture in a multitasking environment, on the other hand, adds a margin of safety to a shared-resource bus, such as Intel's Multibus. Although one of the biggest benefits of such a bus is the sharing of expensive peripherals among several users' programs, an intimidating problem is that a single program gone haywire can easily corrupt the entire system. A two-port memory, properly configured, circumvents this occurrence. Because each port has its own address, data, and control lines, problems on one side are confined by hardware to that side.

Port of call

As a general rule for multitasking architectures, one port of a two-port memory operates in a local environment, and the other port runs remotely, off the expandable shared-resource bus. The local processor is likely to require a synchronous port to reap the benefit of higher performance. Remote buses, in contrast, are usually configured asynchronously. Unless programmed other-

Dynamic-RAM controllers get in step

Synchronous and asynchronous signals have different requirements for interfacing with a controller. The terms synchronous and asynchronous are conventionally applied to dynamic random-access memory depending on whether it exists in a local or a remote environment, respectively. However, they more properly characterize the dynamic-RAM controllers, for the RAMs themselves need no clocks—the only restrictions as to the start of a memory access cycle involve ensuring that the refresh and precharge requirements are satisfied.

Because the controller decides both when to refresh and whether or not precharge and other timing requirements have been met, it does need a clock. Incoming commands can either always arrive with a fixed relationship to the controller's clock or have no particular relationship to it. The former are, of course, synchronous operations, the latter asynchronous.

The major difference between an asynchronous and a synchronous controller (or port of a controller, in the case of the dual-port 8207) is that the asynchronous controller must first synchronize the incoming commands to its own

internal clock. From that point on, the asynchronous controller looks just like a synchronous device.

Whereas various techniques for synchronization are available off chip, on-chip synchronization is restricted to the resolution and sampling of states of a flip-flop. The incoming command is clocked into a resolving flip-flop. After a predetermined time, a sampling flip-flop reads the state of the resolving flip-flop, thereby synchronizing the command. Assuming that both flip-flops are triggered on the same edge of the controller's internal clock, the fastest that an asynchronous signal can be synchronized is one clock period. The slowest synchronization takes two clock periods; on the average, getting the signals in step takes one and a half clock cycles.

Because the processor typically requires four or fewer clock periods to complete a cycle, adding a cycle and a half for synchronizing increases the access time by approximately 25%. Synchronous controllers are therefore always preferred when the environment permits them, and local environments, such as single-board computers, generally do so.

wise, the 8207 configures one port synchronously, and the other asynchronously. For specific applications, both ports may be programmed as either synchronous or asynchronous (see "Dynamic-RAM controllers get in step," p. 129).

Whether the ports are programmed for synchronous or asynchronous operation, some mechanism must decide which processor will gain access to memory when both request it almost simultaneously. That mechanism consists of arbitration logic that controls access and always leaves one port selected. When a port is selected, its associated control and interface signals are passed directly to the RAM timing logic by the command multiplexer (Fig. 2). Both ports' command and control lines, after being synchronized, go into both the command multiplexer and the arbitration logic.

However, the arbitration logic enables the command multiplexer to pass only commands that appear at the selected port. At the same time as a command appears at a selected port, arbitration logic initiates the cycle-control logic that completes the timing of the RAM cycle that ensues. If a command appears on the unselected port, it will not get through the multiplexer to initiate a RAM cycle but will instead wait in the status-command decoder until the current command is completed, at which time the command multiplexer switches to the unselected port. The arbitration logic will then service this queued access request by starting a new cycle.

The arbitration logic examines all port requests, including the internal refresh port. The refresh-request port is subject to arbitration like the other two ports, except that it is always assigned a higher priority than an unselected external access port. Thus, refreshing can be delayed, at most, one RAM cycle.

While the current RAM cycle is running, the arbiter determines the next cycle to be initiated. Thus, the arbitration time of two or more simultaneous port requests is hidden by the memory cycle time. In other words, in cases where both a selected and an unselected port request access simultaneously, the arbitration time for the unselected port does not extend that port's access time, which is delayed by one memory cycle anyway. Only when an unselected port requests a free memory does the arbitration time slow access, because then the command must pass through the arbitration logic before a RAM cycle can be initiated. To minimize such delays in most cases, there are two arbitration algorithms to be selected by the user.

The first algorithm, intended for multiprocessing environments, automatically returns the arbiter to a designated preferred port, generally the higher-performance, synchronous port. Thus any command on the selected port generally has immediate access, whereas any command arriving at the unselected port must wait.

The second, or last-accessed-port, algorithm, which is applicable in multitasking environments, leaves the most recently accessed port as the selected port. This algorithm optimizes port selection for task passing in a multitasking environment. In task passing, the host processor sends a task to an execution processor; until the task is received, the execution processor seldom accesses memory. Conversely, once the task is passed, the host

processor seldom accesses memory until the task is completed. Thus, the ports are used in spurts.

Because timely refreshing is needed to preserve dynamic-RAM data, a refresh request is always serviced on the next available cycle. The refresh algorithm, however, may be selected by the user. The options available are: no refresh, user-generated single refresh, automatic refresh, or user-generated burst refresh.

No refresh would be selected for applications like bit-mapped-video displays, where continuous, sequential access of all RAM locations itself refreshes every cell periodically. User-generated refresh modes allow the designer greater control over power dissipation, for example, in large memory systems. Automatic refreshing, in which the controller itself times the refresh interval and initiates the operation, lets the designer ignore the refresh requirements entirely. As mentioned, the refresh requests are subject to arbitration just like other access requests. However, once a burst refresh is selected, it remains active until completed.

Cleaning up errors

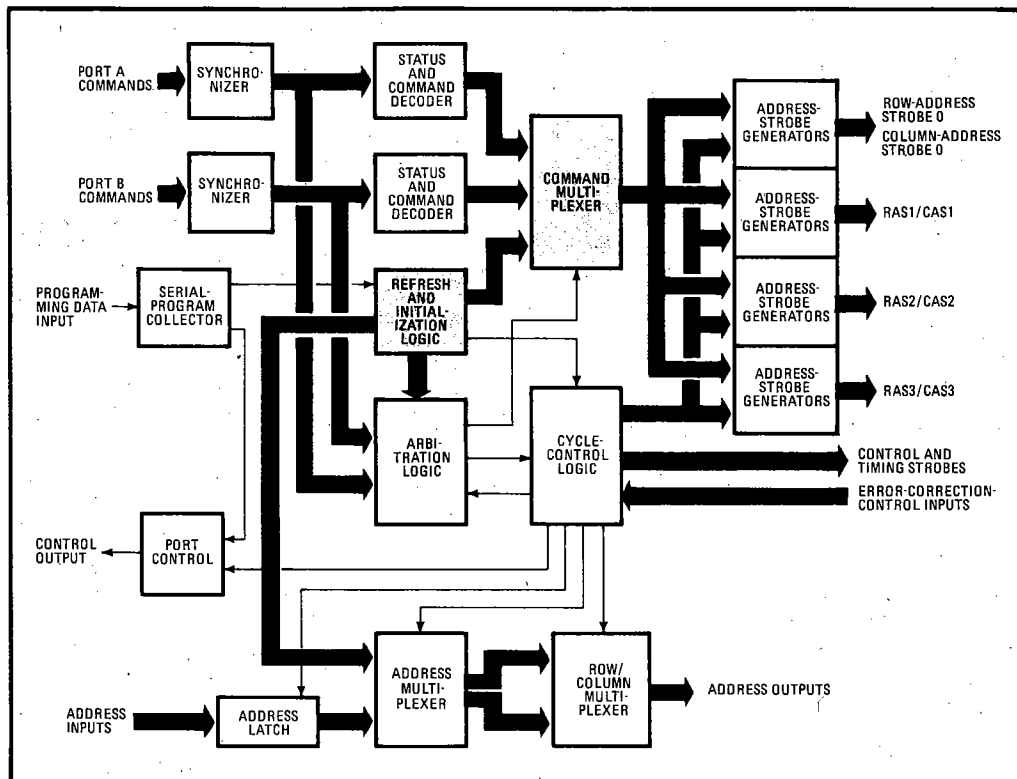
Ensuring data integrity is a major concern in large dynamic-RAM systems, particularly because of their susceptibility to soft errors caused by alpha-particle radiation. Various parity encoding techniques have been developed to detect and correct memory-word errors [*Electronics*, June 2, 1982, p. 153]. The parity bits, called check bits when used for correction as well as detection, are stored in the memory array along with their associated data word. When the data is read, the check bits are regenerated and compared with the stored check bits. If an error exists, whether in the retrieved check bits or in the retrieved data word, the result of the comparison—called the syndrome—gives the location in the group of the bit in error.

Two drawbacks surface in the design of any memory system that is to be protected by error-correction circuitry. First, the memory-word width must be increased to store the check bits; second, extra time must be allotted for the error-correction circuitry to generate the check bits on write cycles, plus more time to regenerate and compare the check bits on read cycles. The 8207 provides several ways to minimize both problems.

Error-correction schemes require a smaller proportion of check bits to protect wider memory words. For example, an 8-bit word needs 5 check bits, for a 63% increase in memory. Put the other way around, 38% of the available memory would be dedicated to the check bits. Six check bits are required to protect a 16-bit data word—only a 27% overhead. Clearly, the wider the memory array, the more economical the error correction.

The 38% overhead necessary to protect such 8-bit-bus machines as the 8088 or 8085 makes error correction an unattractive proposition. However, if the memory width could be doubled, with the 8088 accessing only half a word at a time, the overhead would drop to 27%.

Reading a double-width word, checking for soft errors, and then sending the desired portion of the word to the processor presents no major problems, unlike writing to such an array. The check bits cannot be calculated from only a portion of the word—they must be calculated for



2. Arbitrator's labor. Two external ports plus the internal refresh port can request access to the memory system at once. Arbitration logic decides which to service, based on programmable algorithms. High-speed logic design cuts the delay from input to output switching to 55 ns.

the entire word at once. Whenever the processor writes a partial word to memory, it must first read the entire word, check it, substitute for that portion of the word to be rewritten, and recalculate the check bits. Only then can the entire word be written to memory. The 8207, working in conjunction with the 8206 error-checking and -correction unit, contains mechanisms to expedite this potentially arduous process.

Whenever the 8207 performs a partial-write cycle, it initiates a read-modify-write cycle wherein the entire memory word is first read and latched into the 8206 (Fig. 3). After the retrieved data has been verified as correct, new data is supplied to the RAM, half from the processor and half from the 8206, which also generates the check bits for the entire new word.

Control signals—called byte marks—specify which portion of the new data word is coming from the processor and which from the 8206. The byte marks determine whether the processor or the 8206 drives the RAM data bus—for example, if the 8206 is driving one portion of the data bus, the processor is prevented from driving the same portion. The byte-mark signals simply disable the appropriate transceivers. If, on the other hand, the processor is driving a portion of the RAM data bus, the byte marks change the 8206 data outputs to inputs, allowing

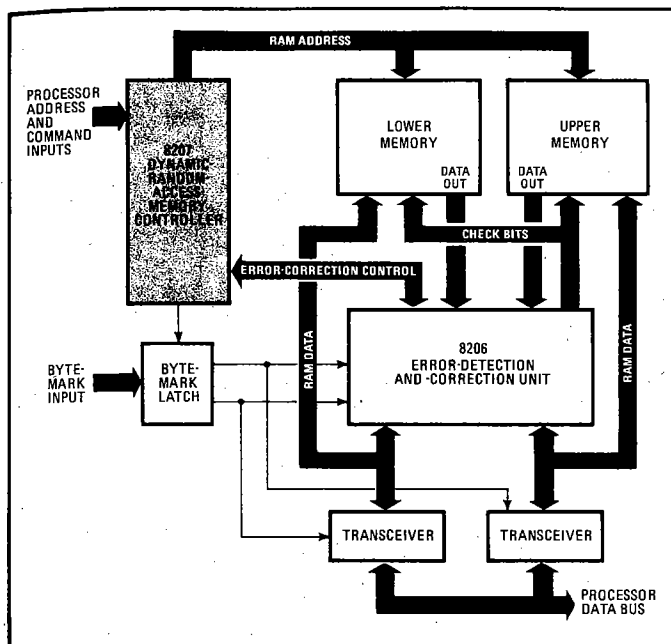
the 8206 to read the data from the processor and calculate new check bits.

The ability of the 8207 to handle memories organized as one, two, or four banks allows tradeoffs between the cost and performance of an error-correction system. For maximum performance, memory would be organized in four banks, each 16 bits wide. In applications requiring error correction, but where maximum performance is not critical, concatenation of RAM banks into two banks of 32-bit words, or even one bank of 64-bit words, can make error correction very economical.

Holding to high performance

Even though the cost of error correction has thus been reduced to where it becomes an attractive solution, the problem remains of minimizing performance degradation. Tackling that challenge depends on the particulars of the configuration, such as whether the memory is to be used with a high-performance local processor, as system memory on a shared-resource bus, or is to be shared between a local high-performance processor and a shared-resource bus.

The method chosen to handle errors depends on the type of bus. Intel's Multibus is the kind that requires data to be valid prior to the issuance of a transfer-



3. Teamwork. The 8206 error-correction chip joins forces with the random-access-memory controller so that an 8-bit-bus processor may utilize the 16-bit-wide memory that is more economical for error-correction schemes. Byte marks configure the data buses for partial-word transfers.

acknowledge signal, in contrast to the local buses of the iAPX-86, -186, and -286 processors. A local bus will usually be synchronous, with a single processor or coprocessor group attached to it; the processor characteristics are known, as is the processor's response to a transfer-acknowledge signal.

With Multibus and other shared-resource buses, the processor types that will eventually be connected are not known in advance, and the buses themselves are generally asynchronous. Hence the time between the transfer-acknowledge signal and data becoming valid is not known. Therefore, the rule with such buses is to acknowledge a transfer only when data is valid. (On some asynchronous buses, the acknowledgment is issued earlier to compensate for synchronization delay at the receiving processor.)

Two basic configurations for checking and correcting errors derive from these system considerations and the fact that it takes longer to correct data than to detect an error. One is for buses that connect to processors and coprocessors receiving a transfer acknowledge prior to data becoming valid, and the other for buses that connect to processors receiving a transfer acknowledge after data is valid. Both configurations are supported by the 8206-8207 team.

Buses among the former type of processors always get corrected data from the 8206, whether an error exists or not, and will carry a transfer acknowledge from the 8207 before data becomes valid on the bus. Though this means data is delayed for error correction on every transaction, the extra delay is immaterial, since it is hidden behind the processor's response time to the transfer-acknowledge signal. By the time the processor requires data, it is

already corrected and on the bus. As a result, system performance is not degraded at all because of single-bit errors.

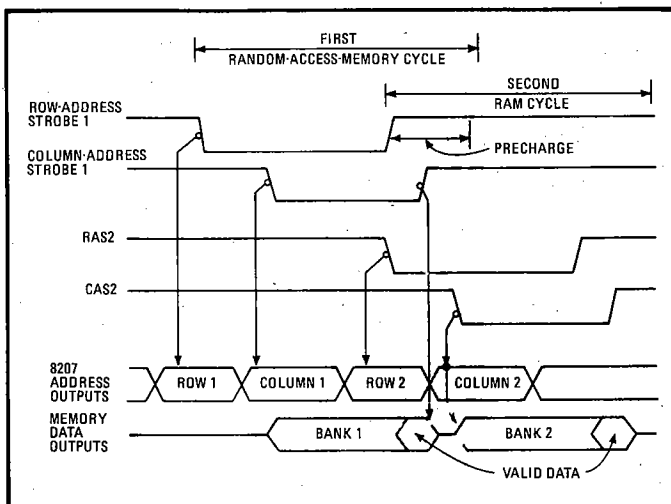
For buses among processors that receive the transfer acknowledge after the data is valid, the 8206 always checks for errors but does not routinely correct data. In this mode, RAM data passes through faster, because the 8207 will issue an acknowledgment sooner. If, however, an error is found, the 8207 will lengthen the cycle, command the 8206 to correct the data, and delay the transfer-acknowledge signal until the corrected data can be placed on the bus. For those buses with an acknowledge-synchronization delay, the 8207 can be programmed to issue the acknowledgment earlier to compensate for the delay.

Power-up problems

Another problem with memories protected by ECC circuits crops up when the power is turned on. At power-up, the data stored in memory is completely random; any attempt to read or perform a partial write will be aborted because the check bits will indicate multiple, and therefore uncorrectable, errors. For processors whose word width is the same as that of the memory array, the processor could simply initialize the entire memory array, taking some additional time and software. For memories whose word width is greater than that of the processor, however, initialization of the memory is not possible unless the error-checking or -correction circuitry is disabled by hardware, for example, by gating off the error flags.

The 8207 is equipped to deal with the initialization problem by itself. At system reset, the 8207 performs

4. Interleaving. Overlapping accesses to different banks increases memory throughput. Once the column-address hold time is satisfied, the 8207 starts a second cycle, pulling the second row-address strobe low.



eight cycles on all banks at once to warm up the dynamic RAMs, a typical RAM requirement for stable operation. The chip then individually initializes all memory locations to 0, adding the proper check bits. Though all memory banks could be initialized in parallel, that would require more power than any other memory operation, calling for a heavier and more expensive power supply needed only at system reset.

One final problem associated with memories protected by error-correction circuitry stems from the fact that only data that is accessed by the processor is corrected. If the processor continually accesses one particular segment of memory, the rest of the array may be accumulating soft errors. The possibility of two soft errors accumulating in a word of seldom accessed memory now becomes significant—and not all double-bit errors are correctable in simple ECC schemes. The 8207 scrubs memories to clean up this problem. During each refresh cycle, one word of memory is read, checked for errors, and if necessary, corrected before data is written back to memory. Because scrubbing occurs during refresh cycles with a read cycle replacing a row-address-strobe-only refresh cycle, no performance penalty is incurred. Scrubbing rids the entire memory of errors at least once every 16 seconds, reducing the probability of two soft errors accumulating in the same word almost to nil.

Bells and whistles

All dynamic RAMs require a recovery period for precharging internal lines after each access. If the processor were immediately to reaccess the RAM, the controller would have to delay it until the precharge time was over. By automatically organizing memory into banks so that sequential addresses are in different banks, the 8207 is usually able to hide the precharge time of one bank behind the access time of another. That organization follows from using the 2 least significant bits of the address to select the bank. Of course, a break in the program flow, such as would be caused by a jump or call

instruction, raises the probability that the same bank may be immediately re-accessed. This probability is less in four-bank memories than in two-bank configurations.

Further performance advantages are gleaned by organizing memory into multiple banks. For example, the 8207 can speed throughput by pipelining cycles. Once the row and column addresses to one bank have been latched, the controller sends the row address for the next cycle to the next bank (Fig. 4).

The 8207's manifold features can be tailored to a given system with the use of a serial programming pin. This pin can either be strapped high or low to select one of two default modes or be programmed by means of a shift register. The external register is completely controlled by the 8207, eliminating any local processor support. Sixteen bits are shifted into the 8207 to configure up to nine different features. The bits are arranged in order of increasing importance; using a shift register with less than 16 bits permits just those features needed to be programmed.

Programmable features of the processor interface include the choice of arbitration algorithm, clock compensation, and preferred port. At the RAM interface, the user can specify fast or slow memory chips, indicate bank configuration, and select the optimal refreshing scheme. In anticipation of the next generation of 256-K dynamic RAMs, the 8207 can support a 256-row-1-millisecond refresh convention, in addition to the 128-row-2-ms one for current 16- and 64-K parts.

Helping facilitate system design is a self-programming processor interface. By decoding the command input pins at power-up, the 8207 automatically determines whether it is connected to the status lines of an 8086, iAPX-286 or to the command lines of the Multibus. Because the 8207 can directly decode the status lines of Intel microprocessors, it can anticipate the next memory cycle and start a new cycle before actually receiving a command. This extra pipelining enables the designer to specify slower RAMs than would otherwise be required. □